

PERSON TAGGING IN AN IMAGE PROCESSING SYSTEM  
UTILIZING A STATISTICAL MODEL BASED ON BOTH APPEARANCE AND GEOMETRIC  
FEATURES

5 Field of the Invention

The present invention relates generally to the field of image processing, and more particularly to techniques for implementing a person tagging feature in an image processing system.

10 Background of the Invention

Person tagging in image processing systems refers generally to the process of characterizing a person observed in an image or sequence of images of a video signal, and using the characterization to determine if the same person is present in one or more subsequent  
15 images. A detected person is "tagged" by association with the characterization, and can thereby be identified as the tagged person in subsequent images. The process of person tagging is thus distinct from a person recognition process in that it does not necessarily involve definitive identification of a given person as being a  
20 particular known individual. Instead, it simply generates an indication that a person in a current image is believed to match a person detected in a previous image. The person tagging process is also referred to as person matching.

Conventional person tagging generally involves the use of either  
25 appearance-based or geometry-based detection algorithms. The appearance-based algorithms include techniques such as template matching and color histograms. Examples of features used in geometry-based algorithms include size, shape, etc. The conventional techniques, however, have been unable to combine appearance and  
30 geometric features in a manner which provides more efficient and effective person tagging for an image processing system.

Summary of the Invention

The present invention solves the above-noted problem of

conventional person tagging techniques by providing a method and apparatus in which appearance features and geometric features are both incorporated into a statistical model of a particular tagged person. The statistical models generated for a given set of persons present in images of a given video segment or other image sequence may be used for detection, location and tracking of the persons in subsequently-processed images.

In accordance with one aspect of the invention, an image processing system processes a sequence of images to generate a statistical model for each of a number of different persons to be tagged so as to be identifiable in subsequent images. The statistical model for a given tagged person incorporates at least one appearance feature, such as color, texture, etc., and at least one geometric feature, such as shape or position of a designated region of similar appearance within one or more images. The models are applied to subsequent images in order to perform a person detection, person location and/or person tracking operation. An action of the image processing system is controlled based on a result of the operation.

In accordance with another aspect of the invention, the statistical model for a given tagged person may be generated by separating one or more images into a number  $N$  of different regions of similar appearance.

In accordance with a further aspect of the invention, the statistical model generated for a given person may be in the form of a likelihood probability function which indicates the likelihood that the person is present in a given image or set of images.

As noted previously, a significant advantage of the present invention is that it utilizes statistical models which incorporate both appearance and geometric features. The use of models which combine these different types of features significantly improves the performance of the person tagging process. For example, such an approach ensures that the system will be less likely to confuse persons crossing one another or persons partially occluded by other

objects in given image sequence.

The present invention can be used in a wide variety of image processing applications, such as video conferencing systems, video surveillance and monitoring systems, and human-machine interfaces.

5

### Brief Description of the Drawings

FIG. 1 is a block diagram of an image processing system in which the present invention may be implemented.

FIG. 2 illustrates an example person tagging process in accordance with the present invention.

FIG. 3 illustrates a translation operation that may be utilized in a person tagging process in accordance with the present invention.

FIG. 4 is a flow diagram of an example person tagging process in accordance with the present invention.

### Detailed Description of the Invention

FIG. 1 shows an image processing system 10 in which person tagging techniques in accordance with the invention may be implemented. The system 10 includes a processor 12, a memory 14, an input/output (I/O) device 15 and a controller 16, all of which are connected to communicate over a set 17 of one or more system buses or other type of interconnections. The system 10 further includes a camera 18 that is coupled to the controller 16 as shown. The camera 18 may be, e.g., a mechanical pan-tilt-zoom (PTZ) camera, a wide-angle electronic zoom camera, or any other suitable type of image capture device. It should therefore be understood that the term "camera" as used herein is intended to include any type of image capture device as well as any configuration of multiple such devices.

The system 10 may be adapted for use in any of a number of different image processing applications, including, e.g., video conferencing, video surveillance, human-machine interfaces, etc. More generally, the system 10 can be used in any application that can benefit from the improved person tagging capabilities provided by the

present invention.

5 In operation, the image processing system 10 generates a video signal or other type of sequence of images of a person 20. The camera 18 may be adjusted such that the person 20 comes within a field of view 22 of the camera 18. A video signal corresponding to a sequence of images generated by the camera 18 is then processed in system 10 using the person tagging techniques of the invention, as will be described in greater detail below. An output of the system may then be adjusted based on the detection of a particular tagged person in a given sequence of images. For example, a video conferencing system, human-machine interface or other type of system application may generate a query or other output or take another type of action based on the detection of a tagged person. Any other type of control of an action of the system may be based at least in part on the detection of a tagged person.

15 Elements or groups of elements of the system 10 may represent corresponding elements of an otherwise conventional desktop or portable computer, as well as portions or combinations of these and other processing devices. Moreover, in other embodiments of the invention, some or all of the functions of the processor 12, memory 14, controller 16 and/or other elements of the system 10 may be combined into a single device. For example, one or more of the elements of system 10 may be implemented as an application specific integrated circuit (ASIC) or circuit card to be incorporated into a computer, television, set-top box or other processing device.

25 The term "processor" as used herein is intended to include a microprocessor, central processing unit (CPU), microcontroller, digital signal processor (DSP) or any other data processing element that may be utilized in a given image processing system. In addition, it should be noted that the memory 14 may represent an electronic memory, an optical or magnetic disk-based memory, a tape-based memory, as well as combinations or portions of these and other types of storage devices.

09703423-110100

The present invention provides improvements over conventional person tagging techniques through the use of statistical models based on both appearance features and geometric features. The term "tagging" as used herein refers generally to the generation of a statistical model characterizing a particular person in one or more images of a given image sequence. A person that has been "tagged" in this manner can then be detected, located and/or tracked in one or more subsequent images of the same sequence or of another sequence.

FIG. 2 illustrates an example of a person tagging process in accordance with the present invention. An image 25 which includes person 20 is generated and processed in system 10 such that the image is segmented into a number  $N$  of different regions of similar appearance. The index  $r$  is used to identify a particular one of the regions. In this example, the image 25 is segmented into a total of  $N = 3$  different regions corresponding to portions 26-1, 26-2 and 26-3 of the original image 25.  $P(I|\Omega)$  denotes the likelihood probability function of a statistical model generated for a given person  $\Omega$ , and indicates the likelihood that the person  $\Omega$  is present in a given image  $I$ . The likelihood probability function  $P(I|\Omega)$  of the statistical model for person  $\Omega$  may be computed as

$$P(I|\Omega) = \sum_{r=1,2,\dots,N} P(R_r|\Omega)P(r|\Omega),$$

where  $R_r$  is a function of at least one appearance feature and at least one geometric feature. The appearance features may include color, texture, etc., and the geometric features may include region shape as well as relative region position within the image.

The general person tagging process illustrated in FIG. 2 involves building statistical models of persons from one or more images and using those models for detection and location of the tagged persons in subsequent images.

The process can also be configured to provide tracking of a tagged person, as will now be described in detail in conjunction with FIG. 3. Let  $P(I | T, \xi, \Omega)$  be the likelihood probability function of the statistical model of the person  $\Omega$ .  $T$  is a linear transformation used to capture global motion of the person in the image space, and  $\xi$  is a discrete variable introduced to capture the state of the local motion of the person at a given point in time, where the term "local motion" is intended to include articulated motion, i.e., the relative motion of different parts of a whole. For example, the position of a person in a room can be obtained from the linear transformation  $T$ , while the pose of the person (standing, sitting, etc.) can be determined from the discrete variable  $\xi$ .

FIG. 3 illustrates the operation of the linear transformation  $T$ . As shown in the figure, the linear transformation  $T$  is used to obtain a sub-window 30 of the image  $I$  that is invariant to rotation and scale. It may be implemented, for example, using a bilinear interpolation technique with a reference point  $x_c$  in the input image  $I$ , a rotation angle  $\theta$ , and a scaling factor  $s$ .

The above-noted local motion is modeled using a discrete set of states  $\{\xi_1, \xi_2, \dots, \xi_M\}$  of the variable  $\xi$  to capture  $M$  different poses of the person  $\Omega$ .

The detection and location of the person  $\Omega$  in the image  $I$  in the person tagging process of the invention may be implemented using the following maximum likelihood search:

$$T^* = \arg \max_T \sum_{\forall \xi} P(I | T, \xi, \Omega) P(\xi | \Omega).$$

Tracking a tagged person, in contrast to detection and location, takes advantage of the history of the known positions and poses of the person from previous images, e.g., previous frames of a given video segment. For a video segment  $V_t = \{I_0, I_1, \dots, I_t\}$ , the likelihood

probability  $P(V_t | T_t, \xi_t, T_{t-1}, \xi_{t-1}, \dots, T_0, \xi_0, \Omega)$  is maximized to obtain the optimum trajectory of the person  $\{T_0, \xi_0, T_1, \xi_1, \dots, T_t, \xi_t\}$ . This maximum likelihood search provides tracking of a tagged person, and can be efficiently implemented using well-known conventional techniques such as the Viterbi algorithm or a forward-backward algorithm.

The likelihood probability of a video sequence can be written in terms of the likelihood probability of individual frames as follows:

$$P(V_t | T_t, \xi_t, T_{t-1}, \xi_{t-1}, \dots, T_0, \xi_0, \Omega) = \\ P(I_t | T_t, \xi_t, \Omega) P(T_t | T_{t-1}, \dots, T_0, \Omega) P(\xi_t | \xi_{t-1}, \dots, \xi_0, \Omega)$$

where  $P(T_t | T_{t-1}, \dots, T_0)$  characterizes global motion model and could be implemented using, e.g., a Kalman filter, and  $P(\xi_t | \xi_{t-1}, \dots, \xi_0, \Omega)$  characterizes local motion, and could be implemented as a first order Markov model using a transition matrix.

In accordance with the invention, different statistical models of the type described above are generated, one for each person present in a given video segment or other type of image sequence. The person tagging process can then provide detection, location and tracking by associating the trajectory of each tagged person with an identifier of the best matching model.

As noted previously, a significant advantage of the present invention is that it utilizes statistical models which incorporate both appearance and geometric features. The use of models which combine these different types of features significantly improves the performance of the person tagging process. For example, it ensures that the system will be less likely to confuse persons crossing one another or persons partially occluded by other objects in the sequence of video frames.

The generation of the statistical models based on both appearance and geometric features will now be described in greater detail. For simplicity and clarity of illustration, the pixels in an image  $I$  of a





determined in accordance with parameters of well-known color spaces such as RGB, HIS, CIE, etc. The texture features may be obtained using well-known conventional techniques such as edge detection, texture gradients, Gabor filters, Tamura feature generation, etc.

5        FIG. 4 is a flow diagram summarizing the above-described person tagging process of the present invention. In step 40, a video segment or other type of image sequence is processed to generate an appearance and geometry based statistical model  $P(I | T, \xi, \Omega)$  for each person  $\Omega$  to be tagged. In step 42, the resulting model or set of models is stored  
10 in a memory of the image processing system, e.g., in memory 14 of system 10. Finally, in step 44, one or more subsequent images are processed using the stored models to perform at least one of a person detection, person location and person tracking operation. The one or more subsequent images may be subsequent images from the same video  
15 segment or other image sequence, or from a different image sequence. The processing operations of steps 40, 42 and 44 may be carried out using software executed by processor 12 of system 10.

00703423-110100  
20        The above-described embodiments of the invention are intended to be illustrative only. For example, the techniques of the invention can be implemented using a variety of different person tagging processes, including processes involving any one or more of person detection, person location and person tracking. In addition, the invention can be used to provide person tagging capability in a wide  
25 variety of applications, including video conferencing systems, video surveillance systems, and other camera-based systems. Furthermore, the invention can be implemented at least in part in the form of one or more software programs which are stored on an electronic, magnetic or optical storage medium and executed by a processing device, e.g., by the processor 12 of system 10. These and numerous other embodiments  
30 within the scope of the following claims will be apparent to those skilled in the art.